(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 825 657 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
**13.10.2010 Bulletin 2010/41**

(21) Application number: **05800938.2**

(22) Date of filing: **15.11.2005**

(51) Int Cl.:
***H04L 29/14*** *(2006.01)* ***G06F 11/20*** *(2006.01)*
***H04L 12/413*** *(2006.01)* ***H04B 1/74*** *(2006.01)*
***H04L 1/22*** *(2006.01)*

(86) International application number:
**PCT/CH2005/000678**

(87) International publication number:
**WO 2006/053459 (26.05.2006 Gazette 2006/21)**

(54) **Reception of redundant and non-redundant frames**

Empfang redundanter und nichtredundanter Rahmen

Réception de trames redondantes et non redondantes

(84) Designated Contracting States:
**AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR**

(30) Priority: **16.11.2004 EP 04405708**

(43) Date of publication of application:
**29.08.2007 Bulletin 2007/35**

(73) Proprietor: **ABB RESEARCH LTD.**
**8050 Zürich (CH)**

(72) Inventors:
- **KIRRMANN, Hubert**
  **CH-5405 Dättwill (CH)**
- **DZUNG, Dacfey**
  **CH-5430 Wettingen (CH)**

(74) Representative: **ABB Patent Attorneys**
**C/o ABB Schweiz AG**
**Intellectual Property (CH-LC/IP)**
**Brown Boveri Strasse 6**
**5400 Baden (CH)**

(56) References cited:
**EP-A- 0 854 610**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 1 825 657 B1

## Description

### FIELD OF THE INVENTION

**[0001]** The invention relates to the field of communication networks in applications such as vehicle control, drive control or substation automation. It departs from a method of receiving frames by a receiving device in a communication network comprising at least two elementary networks or physical links for transmitting redundant frames.

### BACKGROUND OF THE INVENTION

**[0002]** To increase the availability of data networks in vital or otherwise critical applications such as vehicle control, drive control or substation automation, redundancy is required. Basically, the latter may be provided by duplicating the communication lines, e.g. the electrical or optical transmission lines, as well as at least the physical layer of the corresponding protocol stack. While the sending of duplicate information generally does not represent a major obstacle to a straightforward implementation, receiving and handing of duplicated information may require special provisions.

**[0003]** The communication standard IEC 61375, for the purpose of improving availability to enable its use in rail vehicles, specifies two physically independent elementary networks and introduces a frame-by-frame redundancy, in which a sender sends redundant frames simultaneously over the two lines, and a receiver primarily receives the data over one line while observing the other. US Patent Application Publication US 2003/0147377 A1 applies this principle to redundant switched full-duplex Ethernet type communication networks. For the purpose of filtering duplicates, the sender adds an identifier field to the frames before sending identical frames on both elementary networks. A receiver accepts a frame only if its identifier has not been received previously via the other network, i.e. duplicates are rejected at the link layer. The identifier is a continuous frame number or sequence counter value of successive frames that is restarted from zero after a determined count as disclosed e.g. by the IEEE standard 802.2. A given time window delimits the time for acceptance of a frame. In order to maintain the standard Internet Protocol (IP) frame header, the sequence counter is inserted preferably between the payload and the check sequence.

**[0004]** The aforementioned method works well within a closed environment in which all devices are expected to insert a sequence counter in the frames. However, in particular in the area of substation automation, third-party devices and/or mobile devices such as engineering laptops are often connected to one of the two redundant lines only and do not have knowledge of the particular redundancy protocol used by the majority of devices in the network. In this case, the receivers are confronted both with frames that carry a sequence counter and frames that do not carry a sequence counter. In this case, the receivers might either interpret the field comprising the sequence counter and appended to the former frames as an error or ignore it as padding. While the TCP and UDP protocols are able to handle this case, other protocols could actually reject the frame. To avoid this, extensive network and device configuration procedures determining for each device if it does or not obey to the redundancy protocol would be required prior to any data communication. In addition, a general problem of Ethernet frames resides in the fact that in some cases, there is no reliable information about the frame length. Therefore, truncating a frame could result in a correct checksum in some cases, leading to an undetected error.

**[0005]** EP 854 610 discloses an Ethernet redundancy method without switch-over for exchanging measurement data and control commands. The method discards duplicates (i.e. redundant frames) at a low level control section of the receiving device by means of a data identifier, but does not provide for an active detection of singular frames deprived of said data identifier. In detail, singular transmission data addressed to a receiving device is stored in a reception area of a buffer section. In the subsequent step of checking if identical data has already been received over the redundant system's communication line, no such identical data is found, and the singular data is actually forwarded to the respective reception queue as well as to the upper level application. After some time, the abnormality process removes the singular data from the reception queue, as for these no duplicate will be received within a predetermined timeout period. The occurrence of singular data is treated as a scarce case of abnormal behavior, and the method inherently processes the singular data to the upper level application without even considering to identify the singular data at the low level control section of the receiving device.

### DESCRIPTION OF THE INVENTION

**[0006]** It is therefore an objective of the invention to enable, with limited or no configuration effort, a sending device to be connected to one of two lines of a redundant communication network and to send frames to a receiving device connected to both lines and prepared to receive redundant frames sent over both lines. This objective is achieved by a method of receiving frames and a computer program according to the claims 1 and 9. Further preferred embodiments are evident from the dependent patent claims.

**[0007]** According to the invention, the receiving device performs a discrimination or differentiation between non-redundant or singular frames transmitted over one line only, and duplicated or redundant frames transmitted, or at least intended to be transmitted, over both lines. Any frame identified as a singular frame, of which no duplicate exists, is directly forwarded for further processing. Any frame identified as a redundant frame, transmitted ac-

cording to a certain redundancy protocol, is subsequently checked for duplicates having been received beforehand. If the latter is not the case, the redundant frame is identified as a first frame and forwarded for further processing. Otherwise, the redundant frame is identified as a second or late frame and ignored, while nevertheless some information may be derived from its reception. Thus, the receiving device is able to communicate over a communication network with sending devices that implement, as well as with sending devices that do not obey, a particular redundancy protocol.

**[0008]** While the aforementioned discrimination could be performed at any protocol layer of the protocol stack implemented at the receiver, identification and elimination of second frames at the data link layer advantageously saves subsequent protocol layers from executing additional tasks related to duplicate rejection. Nevertheless, as in case of doubt about their redundancy or origin frames are forwarded to a subsequent protocol layer, no frame is disregarded or dropped accidentally. The subsequent protocol layer may be a higher-level layer further inspecting forwarded doubtful frames, or may be a further, singular link layer of a redundancy or relaying device sending forwarded frames over a further, singular line or a non-redundant bus.

**[0009]** A redundancy identifier is added to the two redundant frames as part of the redundancy protocol at the sending device. The redundancy identifier is calculated on the basis of the frame to be transmitted, and compared with a redundancy identifier calculated at the receiving device based on the received frame. This enables a high confidence level to be achieved with a comparatively short redundancy identifier of only a few bits, as compared to fixed and/or pre-determined identifiers serving the same purpose.

**[0010]** Several methods such as checksum or hashing can be used to identify redundant frames, but in a simple embodiment, the value of the redundancy identifier is set equal to the length of the frame. This calculation-effective exploitation of frame-specific information additionally allows the receiver to detect truncated frames with accidentally correct checksums.

**[0011]** In a preferred embodiment of the invention, a redundant frame received over a first line is identified as a first frame of a pair of redundant frames with identical information by means of a frame or sequence counter value transmitted with the received frame. The value of the frame received is compared to a counter or drop window of a predefined maximum length, thus no large tables of used counter values have to be maintained. Although it might be helpful for network management to keep track of the last time when a frame was received from a device, the present method makes no use of a time window.

**[0012]** Advantageously, the aforementioned drop window is delimited by a sync counter value. As soon as the received frames on one of the lines lag behind the sync counter value of the other line, this is interpreted as a loss of synchronisation between the two lines or an at least temporary unavailability of the former line. Hence, a loss of several frames in sequence that can occur e.g. if a switch is reconfiguring itself by executing e.g. a spanning tree protocol, is handled as well. Albeit the algorithm may fail if the delay between the two lines becomes unpredictable, this will not harm transmission, since in case of doubt no frame is rejected and the decision to use the frame is left to the higher protocol layers.

**[0013]** In other words, preferred embodiments of the redundancy method according to the invention comprise one or more of the following advantageous features. a) A support of mixed redundant and singular data traffic, which allows attacking non-critical devices to only one of the two redundant lines without having the interface of a receiving device declare the other line as faulty. b) A use of a long sequence of identifiers, which eliminates the need to create a timer and keep track of time-outs for each single identifier. c) A use of contiguous identifiers, enabling two counter values per source to be updated and out-of-sequence frames to be detected without going through tables of received identifiers. d) An extendability of the redundancy scheme to cover other requirements such as Ethernet security issues as well.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0014]** The subject matter of the invention will be explained in more detail in the following text with reference to preferred exemplary embodiments that are illustrated in the attached schematic drawings, in which:

- Fig.1 shows a bus topology of a communication network,
- Fig.2 shows a switch topology of a communication network,
- Fig.3 depicts the components and function blocks of a device,
- Fig.4 illustrates a redundant frame,
- Fig.5 shows the evolution of two drop windows,
- Fig.6 depicts a virtually redundant network that shares some physical elements,
- Fig.7 shows a redundancy box, and
- Fig.8 illustrates a redundancy switch in another virtually redundant network.

**[0015]** The reference symbols used in the drawings, and their meanings, are listed in summary form in the list of reference symbols. In principle, identical parts are provided with the same reference symbols in the figures.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

**[0016]** Fig.1 shows a bus topology representing e.g. a field-bus according to IEC 61158 and Fig.2 shows a switch topology of a communication network representing e.g. an Ethernet-type local area network (LAN). In both cases, two or more devices 11, 12, 13, 14 are con-

nected by two lines 1a and 1b that are completely separated from each other to avoid a common source of failure. In the present context, the term line does not only describe a point-to-point physical link, but generally refers to one of two redundant paths that can be electrical, optical or based on another technology. As will be apparent from the following detailed description, the invention is independent of the number of connecting elements, hubs or switches and their respective arrangement, e.g. in a loop or tree topology, and is applicable to any type of network and communication media, as well as any number of redundant lines and even only partially, i.e. section-wise, redundant lines.

**[0017]** A sending device or source subscriber equipment 11, 12 sends the same information 100a, 100b over both lines 1a, 1b, and a receiving device or destination subscriber equipment 13, 14 takes the information from whichever line comes first, ignoring the information from the late line but nevertheless registering its reception. If only one line is disturbed, the information will arrive over the other line. This method offers a completely seamless switchover, since there is no need to repeat a piece of information over a previously silent back-up line in case of disruption. It also offers a complete supervision of both redundant lines with normal traffic, ensuring a high coverage and avoiding a dedicated traffic for supervising a back-up line. In Fig.2, devices 30, 31 do not participate in the redundancy protocol and send singular, i.e. non-duplicated information 100s to only one switch 20b.

**[0018]** Fig.3 shows the components and function blocks of a device for sending or receiving redundant frames according to an embodiment of the present invention. In terms of protocol layer hierarchy, this method basically requires the physical layer and the link layer 4 to be duplicated. Accordingly, the sending device 11 comprises two line transceivers 2a and 2b, and two Ethernet controllers 3a and 3b. Processor 10 of the sending device 11 is responsible for providing data packets and preparing a pair of duplicated frames 100a, 100b there from. The frames to be sent are handed through the two Ethernet controllers 3a, 3b to the transceivers 2a, 2b, which in turn do send the frames over the redundant lines 1a, 1b.

**[0019]** At a receiving device 12, 13, the two frames 100a, 100b of a pair are normally almost simultaneously received. Any difference between the arrival times of the two frames may be due to the independent arbitration procedures for accessing the individual lines 1a, 1b, or due to individual delays in the switches 20a, 20b or in other connecting elements not shown in Fig.2. In any case, it can be expected that the order or sequence of the frames sent by one device 11 over one line 1a is maintained, however a later frame 101a sent over one line 1a may be received prior to the reception of a frame 100b that was sent over the other line 1b before said later frame 101a.

**[0020]** The transceivers 2a and 2b are part of the lowest protocol layer of the information transfer process, i.e. the communication, modem or physical layer where the modulation of the digital data occurs. The two Ethernet controllers 3a, 3b establish the connection to the communication media and receive and send the frames from and to corresponding Ethernet controllers in other network devices. They provide the functional and procedural means to transfer data between network devices and to detect and possibly correct errors that may occur in the physical layer. A processor 10 executes, among other tasks, the two link layer protocols or drivers 4a, 4b that are part of the data link layer. The data link layer is hierarchically situated above the physical layer and handles the data exclusively in digital or binary form. A prominent example is Ethernet for local area networks (LAN). Ethernet is a packet-based computer networking technology that defines wiring and signalling for the physical layer, and packet formats and protocols for the media access control (MAC) sub-layer of the data link layer. The data link protocol generally executed in the driver 4a, 4b is often implemented as the IEEE standard 802.2 prescribes, but the Internet Protocol uses a transparent data link protocol.

**[0021]** In conventional non-redundant data transmission, the link layer 4 of the protocol stack is directly connected to a network layer 7 through a link-to-network layer interface 6. The network layer addresses messages and translates logical addresses and names into physical addresses. The network layer also determines the route from the source to the destination device and manages traffic problems, such as switching, routing and controlling the segmentation of data packets. The Internet Protocol (IP) is a widespread network protocol that sends data as blocks referred to as packets or datagrams comprising a header and a data or payload area. Above the network layer, a transport layer not depicted in Fig.3 is situated, which in turn operates e.g. according to the transmission control protocol (TCP) or user datagram protocol (UDP). On top of the transport layer resides an application layer 9.

**[0022]** In principle, the two link layer protocols 4a, 4b could just send both redundant frames to the network layer 7 and frame duplication could be handled by the higher layer protocols in the processors 10. For instance, a transport protocol such as TCP is designed to reject duplicates that can appear in every network due to routing. Other protocols such as UDP or ARP do not detect duplicates, but assume that the application 9 is not affected by reception of duplicates. The latter is true when the network carries sampled values in industrial applications such as IEC 61850, as a reception of duplicate sampled measurement values does not harm the accuracy of the original measurement. The Generic Object Oriented Substation Event (GOOSE) is specified by the substation automation standard IEC 61850 and involves a reduced protocol stack 8 allowing an application 9 to access the link layer of the Ethernet protocol directly, i.e. bypassing the usual network and transport layers. This is most beneficial to the fast communication between de-

vices on the same "subnet", i.e. without relaying at routers or gateways. By way of example, signals for interlocking or data indicative of a status of the substation and sampled at short intervals of e.g. 4 ms, are generally transmitted this way.

**[0023]** However, filtering of duplicates at the link layer 4 instead of handling the redundant frames in the transport layer reduces the computational load of the processor. Furthermore, and in particular in the context of substation automation and the use of the IEC 61850 standard, it is expected that a number of devices obeying to this standard will be connected to a communication network for substation automation without a priori obeying to a common redundancy protocol. Again, eliminating duplicates at the link layer 4 avoids to subject the higher layer protocols 7,8 of these devices to a double traffic.

**[0024]** According to preferred embodiments of the present invention, one or more additional layers are introduced in the protocol stack. In particular, a merge layer 5 formally part of the link layer 4 and hierarchically atop of the two independent link layer drivers 4a, 4b executes the redundancy protocol, and presents towards the high layer protocols 7 or 8 the same link-to-network layer interface 6 as in the non-redundant case. Thus, as long as the rejection protocol works properly, the application layer 9 and the high layer protocols 7 or 8 do not notice the presence or absence of redundancy at all. The merge layer 5 itself behaves as a network layer for the link layer drivers 4a and 4b, and keeps track of the health status of two redundant lines 1a, 1b and therefore can signal disturbances to network management, especially the out-of-sync and stale situations occurring on one of the two lines that otherwise would remain undetected. The main task of the merge layer 5 consists in distinguishing frames that participate in the redundancy protocol from non-redundant or singular frames, e.g. by inspecting a redundancy trailer 200 as detailed below.

**[0025]** Fig.4 shows a typical frame used in the Internet over Ethernet protocol, as could be sent over either line 1a or 1b. The frame starts with a preamble 101, followed by the destination MAC address 102 that identifies the destination, followed by the source MAC address 103 that identifies the sender within the same subnet. The next field 104 is the link layer control, which obeys for instance to the 802.2, 802.1q or another link layer protocol to control data traffic between devices sharing the same medium access. Then comes the payload 105, for instance an Internet packet consisting of an IP header, a TCP header and useful data, or a GOOSE packet in IEC 61850. A frame check sequence FCS 106, e.g. a cyclic redundancy check CRC, computed by the bus controller 3a, 3b concludes the frame.

**[0026]** To the attention of the redundancy protocol, up to four fields of a predetermined bit-length each are inserted as a redundancy trailer 200 between the payload 105 and the frame check sequence 106. The latter comprises a sequence counter 201, a line identifier 202 that indicates over which line the frame has been sent or received, a length field 203 which acts as a redundancy identifier and gives the length of the frame over which the FCS 106 is computed, e.g. from field 101 to field 204 inclusive, and an optional identifier field 204. Obviously, the order of the aforementioned fields in the redundancy trailer 200 is not constrained.

**[0027]** The length field 203 duplicates the information in the link layer header 104 when the latter obeys to the IEEE 802.2 protocols, e.g. for bridge-to-bridge communication. However, Internet frames replace the length field with a protocol type, a potential integrity breach that is overcome by this solution.

**[0028]** The line identifier 202 serves network management purposes and allows to check that the network is correctly configured, in which case it should not be possible that a frame of line 1a appears on line 1b. In some situations, for instance during commissioning or in networks with reduced redundancy, the two redundant frames of a pair might be received over the same line, in which case the receiver segregates the lines by the line identifier field 202, applying the same protocol as if the redundant frames had been received over separate lines. At the same time the receiver signals a configuration error, since this situation bears the danger of a single point of failure. Under normal operation, line identifier 202 may be assigned the hexadecimal coding of "A" or "B".

**[0029]** To distinguish frames 100a, 100b sent according to the redundancy protocol from non-redundant or singular frames 100s, a receiver checks if the field of predetermined length located at position 203 has a redundancy identifier that matches the redundancy identifier computed by the controller of the receiver. In other words, if the content of field 203 does not equal the length of the received frame calculated at the receiver, the receiver assumes that this frame does not participate in the redundancy protocol and forwards it to the subsequent layers. To the same purpose, the optional identifier field 204 may be evaluated as detailed further below.

**[0030]** If one or more of the aforementioned fields 203, 204 indicates a redundant frame, this frame is further inspected by evaluating the sequence number in the field located at position 201. The sequence or frame counter 201 of the latest frame serves the purpose of duplicate rejection as described below. In principle, all operations are executed on Modulo-N arithmetic, N being the largest value of the counter plus one, wherein the counter is expected to have 16 bits. Accordingly, while implementing the following algorithm and referring to differences in counter values, particular care has to be taken.

**[0031]** As mentioned above, each sender keeps a send counter initialised at an arbitrary and preferably random value, inserts the same send counter value into both frames of a pair, at the same time it sets the line identifier 202 and calculates the length field 203. The send counter value is increased for each frame pair the sending device transmits in sequence.

**[0032]** As depicted in Fig.5, a receiver maintains, for each sender identified by its MAC address and possibly

link layer service access point (LSAP), and for both lines, a drop window of $L_W$ successive counter values, wherein the window length $L_W$ is an integer comprised between one and a maximum window length $L_{WMax}$. The drop window is delimited by a sync counter CS and the next expected sequence number CN not yet received. By definition, the latter is not part of the drop window, such that $CN - CS = L_W$.

**[0033]** After initialisation, the receiver accepts the counter value of the first frame it receives from a sender as the sync counter CS of that line. Referring to Fig.5, a frame with a counter value of 23 was the first to be received over line a, followed by four more frames over the same line a. The resulting drop window of length $L_W = 4$ is delimited by the sync counter CS = 23 and the next expected sequence number CN = 27 as depicted on the topmost row indicated "Line A1".

**[0034]** Now, assuming that the receiver receives a frame over line b with a current counter value $CC_B$ of 25, the receiver verifies that the current counter $CC_B$ has the expected value, otherwise an out-of-sequence error will be generated, and compares the current counter $CC_B$ with the drop window of the other line a. As the counter $CC_B$ is within the drop window of the other line, this frame is ignored as it has already been received over line a. The next expected sequence number for line b is incremented to $CN_{B2} = CC_B + 1$. The drop window of line b is reduced to its minimum length by setting the new sync counter $CS_{B2}$ to $CC_B$, and the drop window of line a is shortened correspondingly by setting the sync counter $CS_{A2}$ to $CC_B$ as well, assuming that sequence number will not run backwards and referring to rows "Line A2" and "Line B2".

**[0035]** Next, assuming that the receiver receives, over line a, the next expected frame with a current counter value $CC_A$ of 27. As this value $CC_A$ is outside of the current window for line b, the frame is accepted and passed to the higher levels. The window size $L_{WA}$ is increased from 2 to 3 as the next expected value $CN_{A3}$ is incremented by one and the sync count $CS_{A3}$ is left unchanged, whereas the drop window of line b remains at minimum length as depicted in rows "Line A3" and "Line B3". If a number of frames is received successively over line a, the aforementioned procedure is repeated several times, and the window size $L_{WA}$ is increased up to a maximum $L_{WMax}$. As soon as this maximum is reached, the sync count $CS_A$ is incremented as well, which in turn may make the next expected value $CN_B$ of line b to fall outside of the window of line a. In this case an error counter is incremented that signals that the line b is stalled, and that the lines are out of sync. This out-of-sync state may last until the line b lagging behind has caught up by sending a frame with a sequence number within the drop window of line a.

**[0036]** The optional identifier field 204 may carry additional information, such as a frame checksum computed according to an algorithm agreed upon between sender and receiver, and may serve as an additional identifier for rejecting or further processing of received frames. In particular, repetition of the checksum of the frame as computed by the Ethernet protocol can be used as an additional redundancy identifier. Likewise, the checksum of a frame may be stored at the receiver together with the corresponding counter value, and reject a duplicate frame only when both the counter value and the checksum of the later frame match with a prior redundant frame. The presence or absence of the optional identifier field 204 may be indicated by a bit pattern in field 202.

**[0037]** In Figs.1 and 2 the lines have been depicted as completely separated. In some cases however, and especially during installation, some physical elements of a point-to-point physical link between sender 11 and receiver 13 are shared between the two lines, thus avoiding the need for a complete duplication of the link. For instance, Fig.6 depicts a configuration where the reliability of switch 20 is considered high enough to permit its use in both lines 1a, 1b. Thus, two so-called virtual local area networks 21a, 21b are created instead of two separate physical LANs as in Fig.2. In order to enable the switch 20 to cope with this situation, distinct MAC addresses should be used for the two lines 1a, 1b ending at one and the same device 11, necessitating the Address Resolution Protocol (ARP) to be run independently on both networks.

**[0038]** Fig.7 in turn shows a separate redundancy box 15 that comprises all the elements necessary for relaying frames between lines 1a, 1b and line 1s, i.e. the duplicate transceivers 2a, 2b, controllers 3a, 3b, link layer protocols 4a, 4b, as well as the merge layer protocol 5. The latter is connected to a third singular link layer protocols 4s and controller 3s and a corresponding singular line 1s. The merge layer 5 is responsible for generating redundant frames 100a, 100b for those singular frames received by controller 3s, and for eliminating late redundant frames received by one of the controllers 3a, 3b as described above. Obviously, the functionality of the redundancy box 15 may be incorporated into a redundancy switch 20 as depicted in Fig.8, where basically non-redundant devices 30, 31 are connected to device 11 via two redundant lines 1a, 1b.

**[0039]** The frame or protocol data unit (PDU) according to Fig.4 can advantageously be utilized for the introduction of authentication and encryption to Ethernet in a backward compatible manner as detailed below. Up to now, security protocols have been introduced on the network- (IP-), transport- (TCP-) and higher communication layers, by adding appropriate standardized header fields in their respective PDUs, i.e. the transmitted packet. These fields indicate to the receiving node how to perform message authentication and decryption on the PDU. This presumes that destination nodes are capable of reading those additional security-relevant fields and processing the PDU accordingly. On the Ethernet (link) layer, no security protocol has been standardized to-date. Such Ethernet layer security is however required (i) for protection against link layer spoofing (spoofed link layer source ad-

dress), and (ii) for protection where no higher layer security protocols are available. The latter occurs e.g. in Industrial Ethernets, in particular in the area of substation automation, when the TCP/IP-layers are bypassed for faster transmission within a local area network as specified e.g. by the GOOSE protocol mentioned above. In this case, it is no longer possible to rely on IP- or TCP-layer security protocols for the protection of such "Ethernet-only" PDUs against eavesdropping or tampering.

**[0040]** A straightforward way to introduce additional protocols such as security protocols is by defining appropriate fields in the link layer control (LLC) part 104 of the Ethernet packet header that provide for the necessary trigger: Nodes which are able to recognize these values then know how to interpret further header fields pertaining to the encryption and/or authentication and can process the PDU accordingly. This solution requires that there is a standard specifying those fields. Apart from the Ether type assignments above, no further Ethernet security protocol standard has however been established to-date. But even with such a standard, this solution requires that all nodes in the network must be able to recognize these fields. If a source transmits an Ethernet packet by unicast or multicast carrying security protocol fields, "incompatible" receiving nodes that cannot recognize those fields will not be able to correctly read the rest of the packet. Due to the lack of implementation of standardized Ethernet security protocols in today's Ethernet network interface cards, such a backward compatibility problem becomes often an insurmountable hurdle for the introduction of Ethernet security.

**[0041]** In order to introduce Ethernet security in a manner compatible with Ethernet receiving nodes that are not participating in the security scheme, a similar method as introduced above for the introduction of redundancy information can be used. In analogy to the redundancy identifier that is part of the redundancy trailer 200 appended at the end of the PDU as depicted in Fig.4, a security identifier, preferably in the form of a frame length field, is introduced. To distinguish PDUs according to the security protocol from others, a receiver checks that the value located at the position of the length field 203 is identical to the physical packet length as determined by the controller of the Ethernet receiver. If this is not the case, the receiver assumes that this frame does not participate in the security protocol and forwards it to the upper layers without further processing. If the values agree, the receiver decodes the remaining fields of the redundancy trailer according to the format defined by the Ethernet security protocol and acts accordingly. By way of example, the security fields of interest can be taken from the definition of the known IPSec headers, specifically the encapsulating security payload (ESP) header. Accordingly, said trailer fields comprise (i) a security parameter index pointing to the relevant entry in the security policy database (indicating the encryption and/or authentication algorithm types and keys), (ii) a sequence number for replay protection, (iii) a message authentication code (MAC). The latter may cover the complete PDU for the verification of packet integrity and packet source.

**[0042]** Nodes not understanding the Ethernet security protocol are able to read a non-encrypted packet, albeit without being able to verify integrity and authenticity. Reception without message authentication or without encryption may be sufficient for non-critical receivers, e.g. of multicast packets. Encryption, if enabled in transport mode, must however exclude certain Ethernet addressing fields, in order to allow Ethernet switches to perform address-based switching even if they do not participate in the Ethernet security protocol.

**[0043]** Ethernet receivers not participating in the security protocol will simply not read the appended fields, but are still able to read the PDU correctly and to process the frame check sequence (FCS) for error detection. Thus, this scheme allows the compatible introduction of new Ethernet security protocols. If both a security protocol and a redundancy scheme are employed in the network, a combined redundancy and/or security identifier may be used to indicate this fact, complemented by additional flags signaling the presence of either one or both in the PDU.

LIST OF DESIGNATIONS

**[0044]**

| | |
|---|---|
| 1a, 1b, 1s | lines |
| 2a, 2b, 2s | line transceivers |
| 3a, 3b, 3s | Ethernet controllers |
| 4a, 4b, 4s | link layer protocols / drivers |
| 4 | link layer |
| 5 | merge layer |
| 6 | interface |
| 7 | network and upper layers |
| 8 | GOOSE protocol stack |
| 9 | application |
| 10 | processor |
| 11, 12 | sending devices |
| 13,14 | receiving devices |
| 15 | redundancy box |
| 20, 20a, 20b | switches |
| 21a, 21b | virtual LANs |
| 30, 31 | singular / non redundant devices |
| 100a, 100b | redundant frames |
| 100s | singular frame |
| 101 | preamble |
| 102 | destination MAC address |
| 103 | source MAC address |
| 104 | link layer control |
| 105 | payload, higher protocol data |
| 106 | frame check sequence FCS |
| 200 | redundancy trailer |
| 201 | sequence or frame counter |
| 202 | line identifier |
| 203 | redundancy identifier, frame length field |
| 204 | optional identifier field |

## Claims

1. A method of receiving frames by a device (13, 14, 15, 20) in a communication network, wherein the device (13, 14, 15, 20) has a first and a second network controller (3a, 3b) connected respectively to a first and a second line (1a, 1b) for receiving frames (100a, 100b, 100s) sent by a sending device (11, 12, 30, 31) over the network, comprising,

   - determining if a frame (100a, 100b, 100s) received over one of the first and second line (1a, 1b) is a singular frame (100s) transmitted over the one line (1a, 1b) only or if the frame is a redundant frame of a pair of two redundant frames (100a, 100b) transmitted respectively over the first line (1a) and the second line (1b) according to a redundancy protocol,
   - if the frame is determined to be a redundant frame, determining if the frame is a first frame or a second frame of the pair of two redundant frames (100a, 100b), and
   - forwarding the frame (100a, 100s) for further processing unless the frame is determined to be a second frame (100b),
   the step of determining if the frame is a singular or redundant frame comprises,
   - calculating, by the device (13, 14, 15, 20), a redundancy identifier based on the frame as received, the redundancy identifier comprising a length, a hash, or a checksum of the frame, and
   - comparing this calculated redundancy identifier with a redundancy identifier comprised in a field (203, 204) added to a redundant frame at the sending device as part of the redundancy protocol for redundant frames and absent from a singular frame (100s).

2. The method according to claim 1, **characterized in that** the step of determining if the frame is a singular or a redundant frame is executed at a merge layer (5) of the protocol stack of the device (13, 14, 15), wherein the merge layer (5) receives frames from two link layer drivers (4a, 4b) connected to the network controllers (3a, 3b), and forwards identified singular and first frames (100s, 100a) to a subsequent layer (7, 8, 4s) for further processing.

3. The method according to claim 2, **characterized in that** the merge layer (5) presents a singular link-to-network layer interface (6) towards a higher-level protocol layer (7, 8).

4. The method according to claim 2, **characterized in that** the merge layer (5) forwards the identified singular and first frames (100s, 100a) via a singular controller (3s) to a singular line or to a non-redundant bus (1s).

5. The method according to claim 1, **characterized in that** the redundancy identifier is the frame length (203) of the redundant frame (100a, 100b).

6. The method according to claim 1, **characterized in that** the step of determining if the frame is a first frame or a second frame of the pair of two redundant frames (100a, 100b) comprises

   - comparing a counter value ($CC_B$) comprised in a field (201) added to both redundant frames of the pair as part of the redundancy protocol at the sending device (11), to a drop window ($CS_{A1}$-$CN_{A1}$) maintained at the receiving device (13) for the sending device (11), wherein the drop window comprises a number of successive counter values preceding the next expected sequence number CN not yet received.

7. The method according to claim 6, **characterized in that** separate drop windows are maintained at the receiving device (13) for both lines (1a, 1b) and every sending device (11, 12).

8. The method according to one of claims 1 to 7, **characterized in that** the two network controllers (3a, 3b) are Ethernet controllers and **in that** the communication network serves the purpose of substation automation.

9. A computer program for receiving frames by a receiving device (13, 14) in a communication network, which computer program is loadable in and executable on a data processing unit and which computer program performs, when being executed by a data processing unit, the method according to one of claims 1 to 8.

## Patentansprüche

1. Verfahren zum Empfangen von Rahmen durch eine Vorrichtung (13, 14, 15, 20) in einem Kommunikationsnetz, wobei die Vorrichtung (13, 14, 15, 20) eine erste und eine zweite Netz-Steuereinheit (3a, 3b) aufweist, die mit einer ersten bzw. einer zweiten Leitung (1a, 1b) zum Empfangen von Rahmen (100a, 100b, 100s), die von einer Sendevorrichtung (11, 12, 30, 31) über das Netz gesendet werden, verbunden ist, wobei das Verfahren enthält:

   - Bestimmen, ob ein über die erste oder die zweite Leitung (1a, 1b) empfangener Rahmen ein einzelner Rahmen (100s) ist, der nur über die eine Leitung (1a, 1b) übertragen wird, oder ob der Rahmen ein redundanter Rahmen eines Paars von zwei redundanten Rahmen (100a, 100b) ist, die über die erste Leitung (1a) bzw.

die zweite Leitung (1b) gemäß einem Redundanzprotokoll übertragen werden,
- Bestimmen, ob der Rahmen ein erster Rahmen oder ein zweiter Rahmen des Paars von zwei redundanten Rahmen (100a, 100b) ist, falls bestimmt wird, dass der Rahmen ein redundanter Rahmen ist, und
- Weiterleiten des Rahmens (100a, 100s) für die weitere Verarbeitung, sofern nicht bestimmt wird, dass der Rahmen ein zweiter Rahmen (100b) ist,
wobei der Schritt des Bestimmens, ob der Rahmen ein einzelner oder ein redundanter Rahmen ist, enthält:
- Berechnen eines Redundanz-Identifizierers basierend auf dem Rahmen, wie er empfangen worden ist, durch die Vorrichtung (13, 14, 15, 20), wobei der Redundanz-Identifizierer eine Länge, einen Hash-Code oder eine Prüfsumme des Rahmens enthält, und
- Vergleichen dieses berechneten Redundanz-Identifizierers mit einem Redundanz-Identifizierer, der in einem Feld (203, 204) enthalten ist, das in der Sendevorrichtung als Teil des Redundanzprotokolls für redundante Rahmen zu einem redundanten Rahmen hinzugefügt wird und das bei einem einzelnen Rahmen (100s) nicht vorhanden ist.

2. Verfahren nach Anspruch 1, **dadurch gekennzeichnet, dass** der Schritt des Bestimmens, ob der Rahmen ein einzelner oder ein redundanter Rahmen ist, in einer Mischschicht (5) des Protokollstapels der Vorrichtung (13, 14, 15) ausgeführt wird, wobei die Mischschicht (5) die Rahmen von zwei Verbindungsschicht-Treibern (4a, 4b), die mit den Netz-Steuereinheiten (3a, 3b) verbunden sind, empfängt und die identifizierten einzelnen und ersten Rahmen (100s, 100a) zu einer nachfolgenden Schicht (7, 8, 4s) für die weitere Verarbeitung weiterleitet.

3. Verfahren nach Anspruch 2, **dadurch gekennzeichnet, dass** die Mischschicht (5) eine einzelne Verbindungs-Netzschicht-Schnittstelle (6) zu einer Protokollschicht (7, 8) einer höheren Ebene darstellt.

4. Verfahren nach Anspruch 2, **dadurch gekennzeichnet, dass** die Mischschicht (5) die identifizierten einzelnen und ersten Rahmen (100s, 100a) über eine einzelne Steuereinheit (3s) zu einer einzelnen Leitung oder zu einem nicht redundanten Bus (1s) weiterleitet.

5. Verfahren nach Anspruch 1, **dadurch gekennzeichnet, dass** der Redundanz-Identifizierer die Rahmenlänge (203) des redundanten Rahmens (100a, 100b) ist.

6. Verfahren nach Anspruch 1, **dadurch gekennzeichnet, dass** der Schritt des Bestimmens, ob der Rahmen ein erster Rahmen oder ein zweiter Rahmen des Paars von zwei redundanten Rahmen (100a, 100b) ist, enthält:

- Vergleichen eines Zählerwerts ($CC_B$), der in einem Feld (201) enthalten ist, das in der Sendevorrichtung (11) zu beiden redundanten Rahmen des Paars als Teil des Redundanzprotokolls hinzugefügt wird, mit einem Fallfenster ($CS_{A1}$-$CN_{A1}$), das in der Empfangsvorrichtung (13) für die Sendevorrichtung (11) aufrechterhalten wird, wobei das Fallfenster eine Anzahl aufeinanderfolgender Zählerwerte enthält, die der noch nicht empfangenen nächsten erwarteten laufenden Nummer CN vorausgehen.

7. Verfahren nach Anspruch 6, **dadurch gekennzeichnet, dass** in der Empfangsvorrichtung (13) für beide Leitungen (1a, 1b) und jede Sendevorrichtung (11, 12) separate Fallfenster aufrechterhalten werden.

8. Verfahren nach einem der Ansprüche 1 bis 7, **dadurch gekennzeichnet, dass** die zwei Netz-Steuereinheiten (3a, 3b) Ethernet-Steuereinheiten sind und dass das Kommunikationsnetz dem Zweck der Unterstations-Automatisierung dient.

9. Computer-Programm zum Empfangen von Rahmen durch eine Empfangsvorrichtung (13, 14) in einem Kommunikationsnetz, wobei das Computer-Programm in eine Datenverarbeitungseinheit geladen werden kann und in einer Datenverarbeitungseinheit ausführbar ist, wobei das Computer-Programm das Verfahren nach irgendeinem der Ansprüche 1 bis 8 ausführt, wenn es durch eine Datenverarbeitungseinheit ausgeführt wird.

**Revendications**

1. Procédé pour la réception de trames par un dispositif (13, 14, 15, 20) dans un réseau de communication, le dispositif (13, 14, 15, 20) ayant des premier et second contrôleurs de réseau (3a, 3b) respectivement connectés à des première et seconde lignes (1a, 1b) pour recevoir des trames (100a, 100b, 100s) envoyées par un dispositif émetteur (11, 12, 30, 31) sur le réseau, consistant à :

- déterminer si une trame (100a, 100b, 100s) reçue par l'intermédiaire de l'une des première et seconde lignes (1a, 1b) est une trame singulière (100s) transmise par l'intermédiaire d'une seule ligne (1a, 1b) ou si la trame est une trame redondante d'une paire de deux trames redon-

dantes (100a, 100b) respectivement transmises par l'intermédiaire de la première ligne (1a) et de la seconde ligne (1b) conformément à un protocole de redondance,
- s'il est déterminé que la trame est une trame redondante, déterminer si la trame est une première trame ou une seconde trame de la paire de deux trames redondantes (100a, 100b), et
- réexpédier la trame (100a, 100s) en vue d'un traitement supplémentaire sauf s'il est déterminé que la trame est une seconde trame (100b), l'étape consistant à déterminer si la trame est une trame singulière ou redondante consiste à :
- faire en sorte que le dispositif (13, 14, 15, 20) calcule un identifiant de redondance sur la base de la trame telle qu'elle a été reçue, l'identifiant de redondance comprenant une longueur, un code de hachage, ou un total de contrôle de la trame, et
- comparer cet identifiant de redondance calculé avec un identifiant de redondance contenu dans un champ (203, 204) ajouté à une trame redondante au niveau du dispositif émetteur en tant que partie du protocole de redondance pour des trames redondantes et absent d'une trame singulière (100s).

**2.** Procédé selon la revendication 1, **caractérisé en ce que** l'étape consistant à déterminer si la trame est une trame singulière ou redondante est exécutée au niveau d'une couche de fusionnement (5) de la pile de protocole du dispositif (13, 14, 15), la couche de fusionnement (5) recevant des trames de deux pilotes de couches de liaison (4a, 4b) connectés aux contrôleurs de réseau (3a, 3b), et réexpédiant la trame singulière et la première trame identifiées (100s, 100a) à une couche suivante (7, 8, 4s) en vue d'un traitement supplémentaire.

**3.** Procédé selon la revendication 2, **caractérisé en ce que** la couche de fusionnement (5) présente une interface de couche liaison-à-réseau singulière (6) vis-à-vis d'une couche de protocole de niveau plus élevé (7, 8).

**4.** Procédé selon la revendication 2, **caractérisé en ce que** la couche de fusionnement (5) réexpédie la trame singulière et la première trame identifiées (100s, 100a) par l'intermédiaire d'un contrôleur singulier (3s) vers une ligne singulière ou vers un bus non redondant (1s).

**5.** Procédé selon la revendication 1, **caractérisé en ce que** l'identifiant de redondance est la longueur de trame (203) de la trame redondante (100a, 100b).

**6.** Procédé selon la revendication 1, **caractérisé en ce que** l'étape consistant à déterminer si la trame est une première trame ou une seconde trame de la paire de deux trames redondantes (100a, 100b) consiste à :

- comparer une valeur de compteur ($CC_B$) contenue dans un champ (201) ajouté aux deux trames redondantes de la paire en tant que partie du protocole de redondance au niveau du dispositif émetteur (11), à une fenêtre déroulante ($CS_{A1}$-$CN_{A1}$) maintenue sur le dispositif récepteur (13) pour le dispositif émetteur (11), la fenêtre déroulante comprenant un certain nombre de valeurs de compteur consécutives précédant le numéro de séquence CN attendu suivant et non encore reçu.

**7.** Procédé selon la revendication 6, **caractérisé en ce que** des fenêtres déroulantes séparées sont maintenues sur le dispositif récepteur (13) pour les deux lignes (1a, 1b) et pour chaque dispositif émetteur (11, 12).

**8.** Procédé selon l'une des revendications 1 à 7, **caractérisé en ce que** les deux contrôleurs de réseau (3a, 3b) sont des contrôleurs Ethernet et **en ce que** le réseau de communication est utilisé à des fins d'automatisation de stations secondaires.

**9.** Programme informatique destiné à la réception de trames par un dispositif récepteur (13, 14) dans un réseau de communication, lequel programme informatique peut être chargé et exécuté sur une unité de traitement de données, et lequel programme informatique met en oeuvre, lorsqu'il est exécuté par une unité de traitement de données, le procédé selon l'une des revendications 1 à 8.

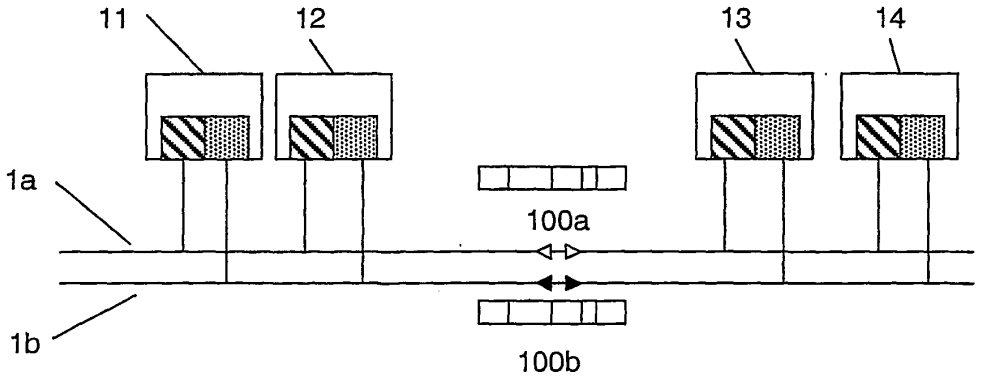

11
12
13
14
1a
100a
1b
100b

Fig. 1

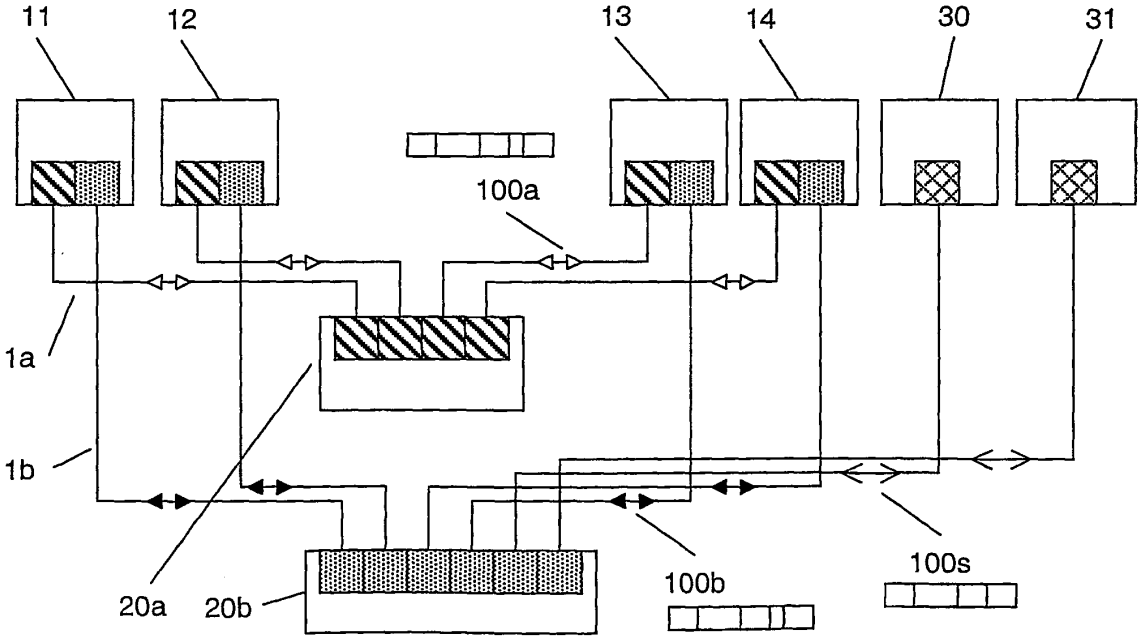

11
12
13
14
30
31
100a
1a
1b
20a
20b
100b
100s

Fig. 2

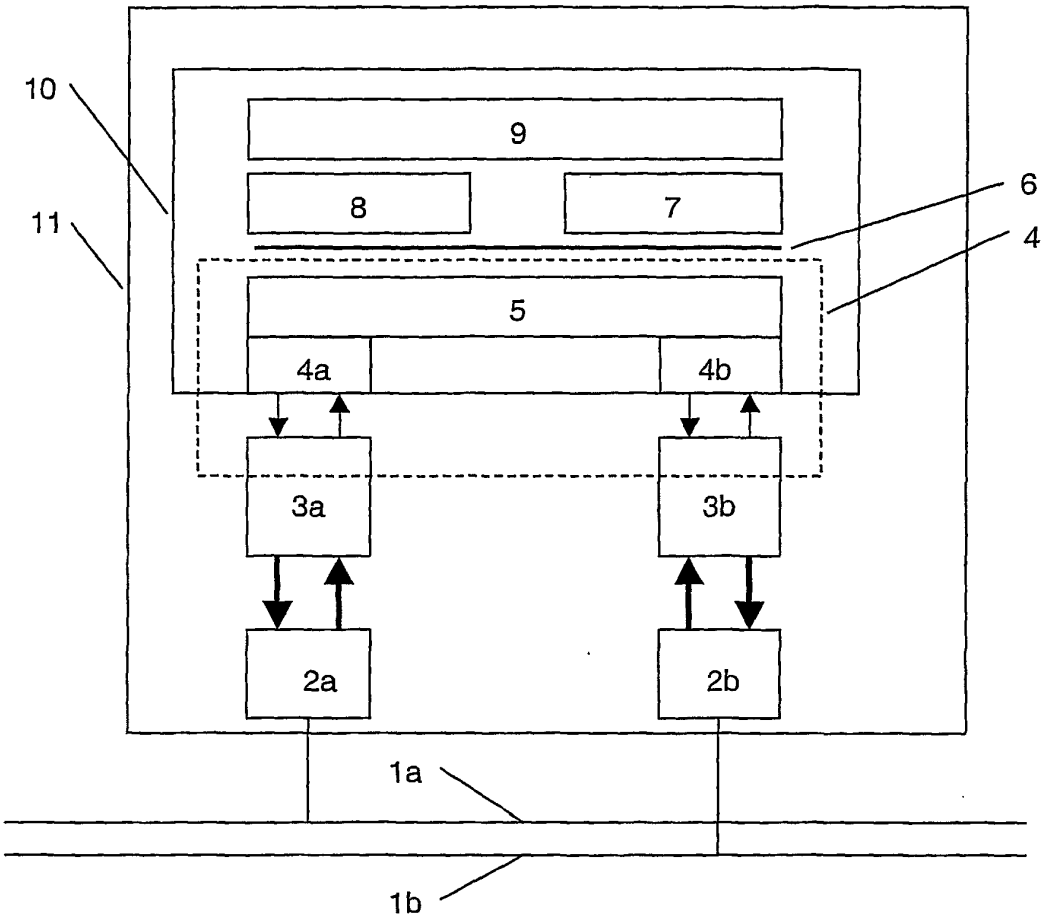

10
11
9
8
7
6
4
5
4a
4b
3a
3b
2a
2b
1a
1b

Fig. 3

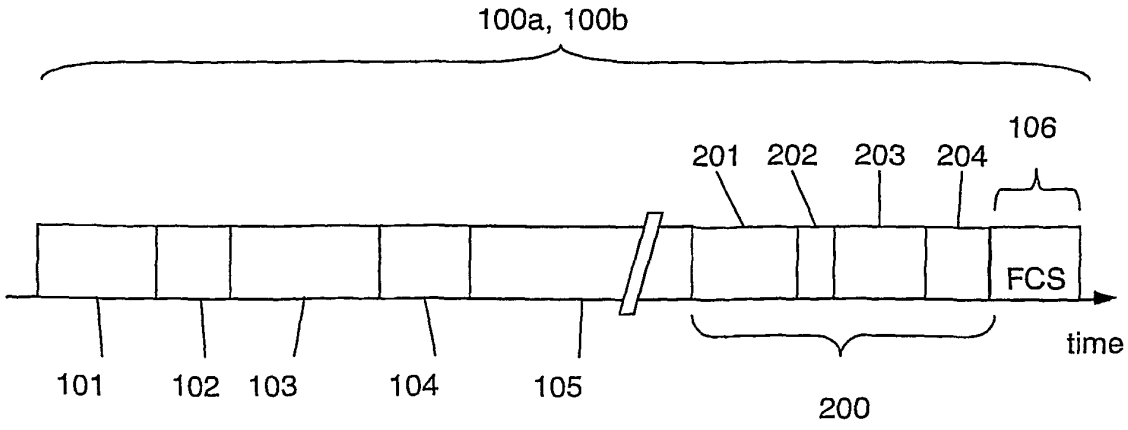

100a, 100b
106
201
202
203
204
FCS
time
101
102
103
104
105
200

Fig. 4

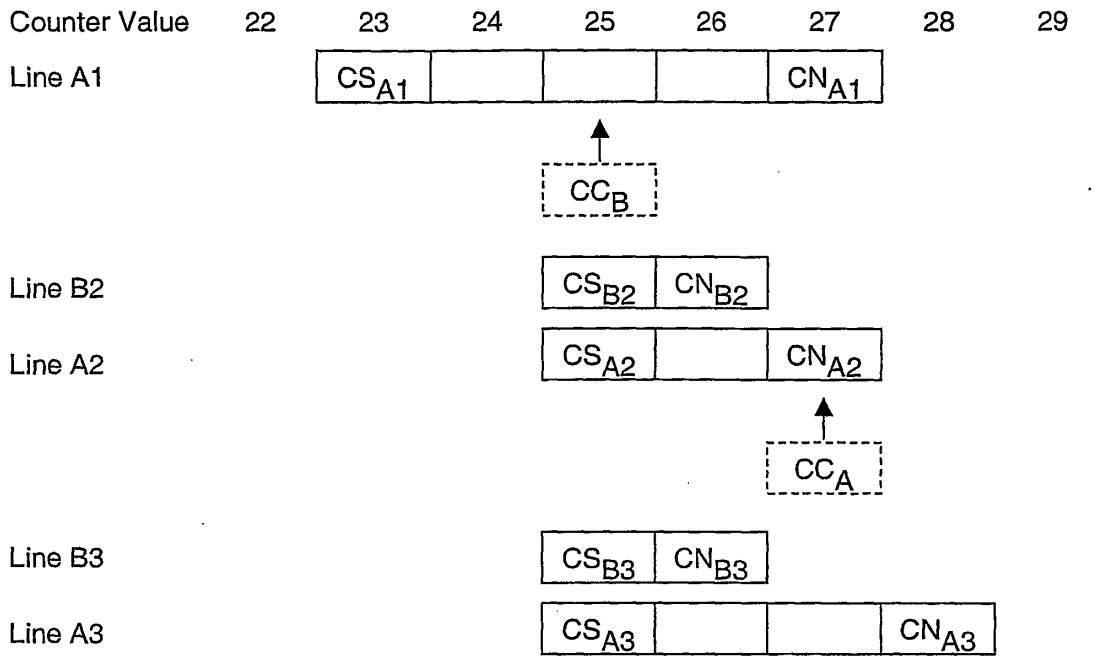

Counter Value
22
23
24
25
26
27
28
29
Line A1
CS_A1
CN_A1
CC_B
Line B2
CS_B2
CN_B2
Line A2
CS_A2
CN_A2
CC_A
Line B3
CS_B3
CN_B3
Line A3
CS_A3
CN_A3

Fig. 5

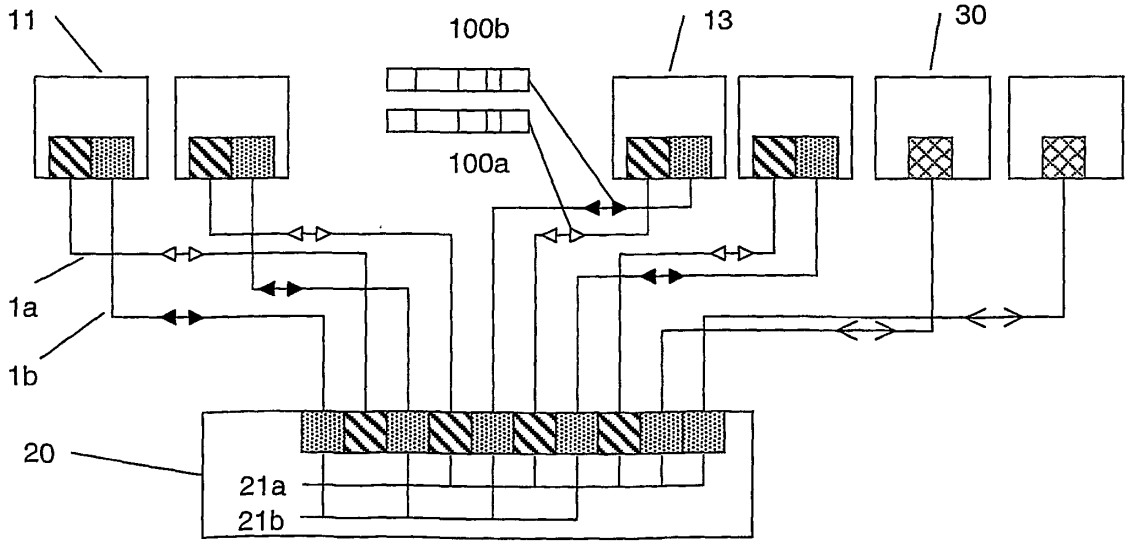

11
100b
13
30
100a
1a
1b
20
21a
21b

Fig. 6

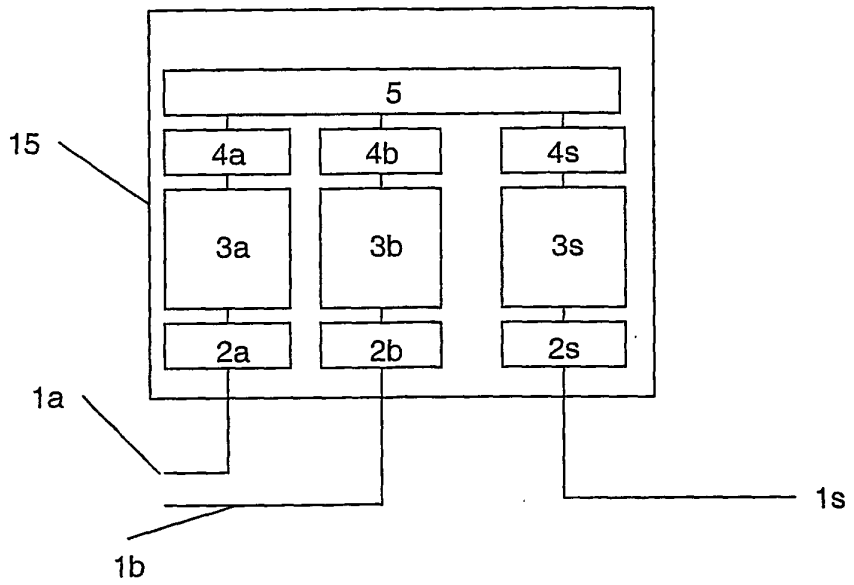

5
15
4a
4b
4s
3a
3b
3s
2a
2b
2s
1a
1s
1b

Fig. 7

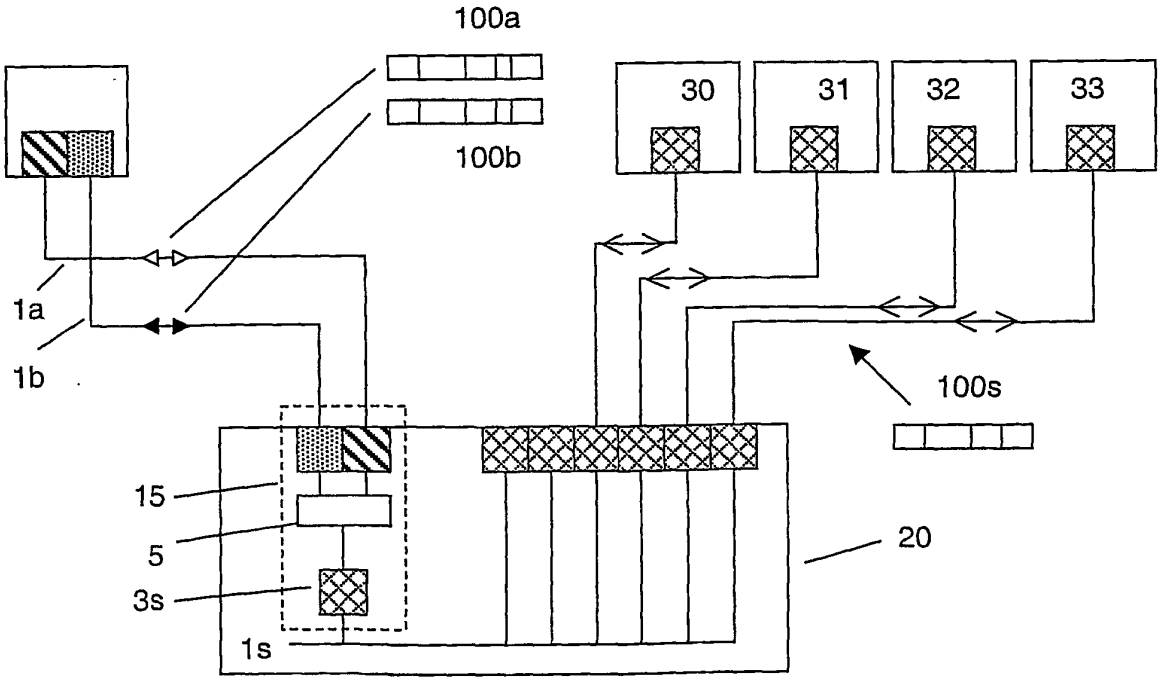

100a
100b
30
31
32
33
1a
1b
100s
15
5
3s
1s
20

Fig. 8

## REFERENCES CITED IN THE DESCRIPTION

*This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.*

### Patent documents cited in the description


- US 20030147377 A1 **[0003]**
- EP 854610 A **[0005]**